

Computational Physics

Peter Hertel

**Physics Department
Osnabrück University
Germany (EU)**

Lectures delivered at
TEDA Applied Physics School
Nankai University, Tianjin, PRC

Numerical, or computational physics is a branch of physics with an age-old tradition. If you can't produce a number, you have achieved nothing says Feynman, and right he is. More often than not this is possible only by computing. This series of lectures introduces standard methods of computational physics. The computational power which is available to the average student or researcher has grown by a factor of thousand or so, within the last twenty years. And with it the availability of high quality software. Therefore, these lectures concentrate on how to use software, and not how to develop it. Quite naturally, this boils down to MATLAB.

Autumn 2010

Contents

1	Introduction	3
2	Matrices	4
3	Numbers	6
4	Functions	8
5	Ordinary differential equations	11
6	Fast Fourier transform	15
7	Fitting data	18
8	Simulated Annealing	23
9	Finite difference method	27
10	Propagation	31
A	Figures Makefile	36
	List of Figures	38

1 Introduction

This series of lectures is about standard numerical techniques for computing in physics. Although teachers prefer examples which can be solved by pencil and paper, or by chalk on the blackboard, most problems cannot be solved analytically. They require computations. Computational physics has a long tradition. However, computational physics cannot be taught any more in the traditional way since the personal computer has become part of everyday life.

In 1980, the first IBM personal computer could address 640 kB of memory, its mass storage was a floppy disk of 360 kB, and it ran at 8 MHz with an 8 bit processor. With 15,000 EUR it was cheap as compared to mainframes. Today, a modest computer, such as my laptop, costs 600 EUR. Its double core 32 bit processor runs at 1.6 GHz, is supported by 2 GB fast storage and a 120 GB hard disk. This amounts to a many-hundred-fold increase in computational power within the last three decades.

Moreover, high quality software to be run on these cheap and highly efficient computers has become cheap. Just study the price list of the MathWorks corporation for MATLAB.

Today, a course on computational physics, in particular an introduction, should concentrate on how to use available and cheap high quality software, and not how to develop it.

I have decided on using MATLAB for various reasons.

First, because the program package is relatively cheap.

Second, because it comes with an integrated development environment such as a command window, a language sensitive editor, a debugger, a workspace inspector, and so forth.

Third, because it includes the software treasures of the last 50 years. In particular, the LINPACK and EISPACK packages¹ of linear algebra, forming the backbone of computational physics, are the core of MATLAB. In fact, MATLAB was conceived as a meta-language for how to use such FORTRAN code without indulging into unnecessary details.

Fourth, MATLAB provides for sparse matrix technology which is essential for partial differential equations.

Fifth, MATLAB makes you think in terms of objects, it is not a procedural language. After a short while, you will think of x as a variable, not of an array of equally spaced representative values.

And last but not least, MATLAB is a powerful graphics machine for visual-

¹Steve Moler is one of the authors of the the LINPACK and EISPACK packages documentation, and one of its main contributors. Today he is the chief scientist of MathWorks. See <http://www.mathworks.com/company/aboutus/founders/clevemoler.html>

izing data, both two and three dimensional. In particular, MATLAB allows for the export of pictures in encapsulated postscript format.

In these lectures on Computational Physics we cannot cover a number of other important fields, such as data acquisition, graphical user interfaces or fine points of graphics programming. Instead, we will restrict ourselves to standard techniques of doing analysis on a computer.

This text addresses students of physics who are used to learn by studying carefully chosen examples. It is neither systematic nor complete.

2 Matrices

The basic objects of MATLAB are matrices. A number is a 1×1 matrix. A row vector is $1 \times N$ matrix, a column vector is an $N \times 1$ matrix. A polynomial is represented by the row vector of its coefficients. A string is a row vector of 2 byte numbers which represent characters. And so forth.

The most important operators are square brackets for aggregation. The comma operator separates objects which are to be aggregated horizontally. The semicolon operator separates objects which are to be aggregated vertically. The semicolon operator, at the end of an expression, suppresses echoing.

Here are a few examples:

```
>> A=[1,2,3;4,5,6];
>> B=[A;[7,8,9]];
```

The `size` function takes in a matrix and returns the numbers of rows and columns. Try

```
>> [r,c]=size(A);
```

Note that output is always aggregated horizontally.

By saying²

```
>> M=3; N=5;
>> a=zeros(M,N);
```

you create an $M \times N$ matrix `a` of zeros. Likewise, `ones` creates matrices of ones.

```
>> B=eye(N);
```

²MATLAB distinguishes between lower and upper case

generates a square 5×5 unit matrix. Check it by typing

```
>> B(2,1)
>> B(2,2)
```

The dash operator stands for Hermitian conjugation. Compare

```
>> x=[1,2+3i,4;0,i,2]
```

and

```
>> xc=x'
```

One of the highlights of MATLAB is the colon operator. It stands for 'from:to'. The colon itself stands for all indices. Thus we could have written

```
>> B=[1:3;4:6;7:9];
```

We now extract the first and the third row:

```
>> C=B([1,3],:);
```

Note that matrices are used just as functions. Arguments are in parentheses and separated by commas. The first argument is a vector of column indexes, the second argument a vector of row indexes.

Another often used function for creating vectors is `linspace`. This function has three arguments: the lower value of an interval, the upper value, and the number of representative points.

```
>> x=linspace(0,10,128);
```

produces a row vector of size 1×128 . Its first element is 0, the last is 10, and linear interpolation in-between.

```
>> x=[0:10/127:10];
```

does the same.

Here the `from:step:to` construct is used. If the difference between `to` and `from` is not an integer multiple of `step`, then `to` is an upper bound. If `step` is omitted, a value of 1 is assumed.

From the 10×10 matrix

```
>> x=rand(10,10);
```

of randomly chosen numbers we may extract the sub-matrix

```
>> y=x(1:2:10,1:2:10);
```

of random numbers with odd indices.

Constructs like

```
>> k=[1,4,9:2:17,31];
```

are allowed as well. MATLAB's syntax rules, if strictly applied, would require

```
>> k=[[1],[4],[9:2:17],[31]];
```

They are, however, sensibly relaxed. In many cases the horizontal alignment comma operator may also be omitted,

```
>> k=[1 4 9:2:17 31];
```

works as well. We shall try to avoid this kind of syntax relaxation although you may encounter it in the help system.

3 Numbers

MATLAB is rather strict with numbers. It adheres to the IEEE³ convention on representing real numbers and algebraic operations with them. It is obvious that real numbers must be approximated. The approximation, however, should be reliable and as good as possible.

There is a predefined number `eps`. It is the smallest number such that 1 and $1+\text{eps}/2$ are identical. On my computer `eps=2.2204e-016`. Roughly speaking, the inherent accuracy is 16 digits. This is sufficient for physics where not more than 6 digits are meaningful.

Rounding errors are randomly distributed. With N operations, the error grows proportional to \sqrt{N} . Hence, 10^{20} operations on real numbers are allowed until the statistical error exceeds the 6 digits limit. If your computer runs at 1.0 GHz, approximately 10^{11} seconds are required to perform such a large number of operations. This is more than 1000 years. Put otherwise: rounding errors, unless crucial, will not pose a problem.

Rounding errors become crucial if a result is the difference between two large, but almost equal numbers. Try

```
>> x=1e-15;  
>> y=((1+x)-1)/x
```

³Institute of Electrical and Electronics Engineers

The result is 1.1102, but should be 1.

Assume you want to calculate $f(x) = \sin(x)/x$ for $x \in [0, 2\pi]$. The following code does is:

```
>> x=linspace(0,2*pi,1024);  
>> y=sin(x)./x;
```

Typing

```
>> y(1)
```

produces NaN, not a number. And this should be so. After all, we have divided zero by zero, and the result is neither 1 nor infinite, it is not defined. The IEEE standard of floating point representation, which is realized by all Intel processors and supported by MATLAB software, provides for a bit pattern which is not a valid number. Any operation with NaN results in NaN. NaNs tend to spread out. For example,

```
>> z=y/sum(y);
```

results in a vector filled with NaNs.

A NaN can be detected by the function `isnan`. `isnan(y(1))` returns 1 which stands for true. `isnan(y(2))` results in 0, or false. `isnan(y)` returns an vector the first component of which is 1, the remaining are 0.

Another special number of the IEEE standard is `Inf` which stands for positive infinity. $1/0$ results in `Inf`, $-1/0$ in `-Inf`, just as $\log(0)$. `Inf/Inf` or `Inf-Inf` are ill defined, i. e. NaN. However, `Inf+Inf` gives `Inf`.

The pseudo-numbers `Inf` and `NaN` solve a long-standing problem. In the old days, you could either demand program abort if an ill-defined arithmetic operation was performed, or the result was arbitrary. With the IEEE standard and high quality software the program is not halted, but errors are realized and documented.

Let us return to $f(x) = \sin(x)/x$. This function, by declaring $f(0) = 1$, can be made to be continuous everywhere.

```
>> y=sin(x)./x;  
>> y(1)=1;
```

is a plausible and obvious solution.

However, you have made use of the fact that it is the first component which is exceptional. Much better is the following solution:

```

>> y=ones(x);
>> k=find(x~=0);
>> y(k)=sin(x(k))./x(k);

```

`find` returns a vector of indices for which `x` is not equal to zero.

A third possibility⁴ is old-style procedural programming:

```

1  for k=1:1024
2      if x(k)==0
3          y(k)=1;
4      else
5          y(k)=sin(x(k))/x(k);
6      end
7  end

```

This is not recommended. Loops in MATLAB are slow, they should be avoided by all means.

A fourth possibility would be

```

>> x=linspace(eps,2*pi,1024);

```

Although it works, this is a dirty trick. The problem is avoided, not solved.

4 Functions

Functions map one or more real or complex variables into real or complex numbers. Functions are essential for physics, they describe relationships between measurable quantities.

Let us study an example. The spectral intensity of black body radiation is described by Planck's famous formula

$$S(x) = \frac{15}{\pi^4} \frac{x^3}{e^x - 1} , \quad (1)$$

where x is short for $\hbar\omega/k_{\text{B}}T$.

(1) is good praxis. Computational physics has to deal with numbers, i. e. dimensionless quantities. The spectral density is a probability distribution, a dimensionless quantity. In fact,

$$\int_0^z dx S(x) = \text{Pr}\{\hbar\omega < zk_{\text{B}}T\} . \quad (2)$$




Figure 1: The black body radiation spectral intensity $S(x)$ where $x = \hbar\omega/k_{\text{B}}T$

In MATLAB you describe a function in a separate m-file. The filename without the .m extension is the name of the function. (1) is realized as follows:

```

1  function s=planck(x)
2  % usage: s=planck(x);
3  % s=(15/pi^4)*x.^3./(exp(x)-1)
4  tiny=1e4*eps;
5  kt=find(abs(x)<=tiny);
6  kn=find(abs(x)>tiny);
7  s(kt)=x(kt).^2;
8  s(kn)=x(kn).^3./(exp(x(kn))-1);
9  s=15*s/pi^4;
10 end %function planck

```

For very small x the denominator is the difference between $\exp(x)$ and 1, two close-by large numbers. We therefore replace $x^3/(\exp(x) - 1)$ by x^2 . `kt` is an array of indexes where x is tiny, while `kn` comprises all indexes where x is not tiny. For both index sets the appropriate formula is applied. Although this seems to be complicated, an x vector of 100,000 entries is transformed in less than 50 ms.

⁴Line numbers instead of the `>>` prompt indicate that we have executed a script file, `test.m` in this case

We may now say

```
>> x=linspace(0,10,1024);
>> plot(x,planck(x));
>> print -deps2 planck.eps
```

Fig. 1 shows the result.

The figure has been exported to an Encapsulated Postscript (level 2) file `planck.eps`. It was transformed into `planck.pdf` by the `eps2pdf` program and included into this L^AT_EX document by saying

```
\FIG{planck.pdf}
{The black body radiation spectral intensity  $S(x)$ 
where  $x=\hbar\omega/k_B T$ .}
```

My `\FIG` macro is defined as follows:

```
1  \newcommand{\FIG}[2]{
2  \begin{figure}[!hbt]
3  \begin{center}
4  \begin{minipage}{0.85\textwidth}
5  \centering{\includegraphics[width=95mm]{#1}}
6  \caption{\label{#1}\small{#2}}
7  \end{minipage}
8  \end{center}
9  \end{figure}
10 }
```

You must have imported the `epsfig` package before.

We have explained how to properly define and how to visualize a function. What else can you do?

Let us numerically check whether (1) is indeed a probability distribution. We approximate $x = \infty$ by $x = 20$, say. The following line performs the integration:

```
>> quad(@planck,0,20)
```

The output 0.99999729919447 (with `format long`) is convincing since the quadrature method `quad` has a predefined tolerance of six digits. And 20 is not yet infinity. `quadl` is better. The following call

```
>> quadl(@planck,0,40)
```

results in 1.00000000003192.

The function, here `planck`, is referred to by its handle `@planck`.

```
>> quad('planck',0,20)
```

will work as well. Although there are subtle differences, we shall not discuss them here.

`quad` or `quadl` are in-built functions which work on functions. You may inquire by typing

```
>> help funfun
```

`fminbnd` is such a function function. It finds out the position of the minimum within prescribed bounds. However, we want to know about the maximum position of our planck function. Therefore we must define the negative of the planck probability distribution. For this we do not require a new m-file, because there is the `inline` statement:

```
>> np=inline('-planck(x)');
```

The string is parsed as we would read it.

```
>> xmax=fminbnd(np,0,20)
```

delivers 2.8214.

We now may ask

```
>> elo=quad(@planck,0,xmax);  
>> ehi=quad(@planck,xmax,20);
```

for finding the amount of energy below and above the spectral intensity maximum. The two numbers should add up to 1. Do they?

5 Ordinary differential equations

We talk about a system which is in a certain state. If the system has N degrees of freedom, the state is described by a vector y_1, y_2, \dots, y_N of variables. The state will change in the course of time. In many cases the time development of the state of a system is adequately described by a system of ordinary differential equations,

$$\dot{y}_j = f_j(t, y_1, y_2, \dots, y_N) . \quad (3)$$

Higher than first derivatives can always be removed. After all, the second derivative is the first derivative of the first derivative, and so forth.

Let us investigate one of the oldest physical problems, the motion of planets in the sun's gravitational field.

Recall that the conservation of angular momentum requires the planet to move in a plane. The planet's location $x_1(t), x_2(t)$ obey the following differential equation:

$$m\ddot{x}_j(t) = -\frac{GmM_\odot x_j}{(x_1^2 + x_2^2)^{3/2}} \quad (4)$$

G is the universal gravitational constant, m the planet's and M_\odot the sun's mass. m drops out. By measuring x in units of the astronomical unit⁵ a and t in units of $\tau = \sqrt{a^3/GM_\odot}$ we arrive at⁶

$$\ddot{x}_j(t) = -\frac{x_j}{(x_1^2 + x_2^2)^{3/2}} \quad (5)$$

Let us introduce $(y_1, y_2, y_3, y_4) = (x_1, \dot{x}_1, x_2, \dot{x}_2)$ such that

$$\dot{y}_1 = y_2 \quad (6)$$

$$\dot{y}_2 = -y_1/r^3 \quad (7)$$

$$\dot{y}_3 = y_4 \quad (8)$$

$$\dot{y}_4 = -y_3/r^3 \quad (9)$$

results, where $r = \sqrt{y_1^2 + y_3^2}$.

We describe this set of four ordinary differential equations by the following derivative field:

```
1 function d=kepler(t,y);
2 r=sqrt(y(1).^2+y(3).^2);
3 d=[y(2);-y(1)./r^3;y(4);-y(3)./r^3];
```

At the command window we say

```
>> [t,y]=ode45(@kepler,0:0.1:50,[1,0,0,0.8]);
```

A Runge-Kutta ordinary differential equation solver `ode45` is invoked. It needs a derivative field, a time span, and a start vector. And this is the result of saying

```
>> plot(y(:,1),y(:,3))
>> print -deps2 kepler1.eps
```

Although the trajectories look like ellipses, they shrink more and more. The reason is not physics, but numerics. We substantiate this remark by plotting the energy in Fig. 3

⁵the mean distance between earth and sun, 149.6×10^{11} m

⁶ $2\pi\tau$ is one year




Figure 2: Planetary motion without explicit accuracy control




Figure 3: Total energy vs. time without accuracy control

```
>> Ekin=0.5*(y(:,2).*y(:,2)+y(:,4).*y(:,4));  
>> Epot=-1./sqrt(y(:,1).*y(:,1)+y(:,3).*y(:,3));  
>> plot(t,Ekin+Epot);
```

We try better by setting a lower relative tolerance (the default being 0.001):

```
>> tol=odeset('RelTol',1e-8);  
>> [t,y]=ode45(@kepler,0:0.1:50,[1,0,0,0.8],tol);
```




Figure 4: Planetary motion with accuracy control

Fig. 4 shows what we expect: closed ellipses.

There are many more options for the ordinary differential equation solver, and there are more such solvers. You should ask for details at the help desktop.

6 Fast Fourier transform

There are three kinds of Fourier transforms: finite discrete, infinite discrete, and continuous. The first is defined by

$$G_j = \sum_{k=0}^{N-1} e^{-2\pi i j k / N} g_k = \sum_{k=0}^{N-1} \Omega_{jk} g_k \quad (10)$$

for $j = 0, 1, \dots, N-1$. The Matrix Ω is

$$\Omega_{jk} = \omega^{-j \cdot k} \quad \text{where } \omega = e^{2\pi i / N} \quad (11)$$

ω is the N th root of 1.

Because Ω / \sqrt{N} is unitary,

$$\sum_{k=0}^{N-1} \Omega_{jk} (\Omega^\dagger)_{kl} = N \delta_{j,l} \quad (12)$$

we obtain for the inverse transform the following expression:

$$g_k = \frac{1}{N} \sum_{j=0}^{N-1} e^{+2\pi i j k / N} G_j \quad (13)$$

Only the finite discrete Fourier transform is of interest in computational physics since the real axis has to be approximated by a finite interval, and an interval by a finite set of representative points.

To be specific, we think about a time interval, $t_k = k\tau$ for $k = 0, 1, \dots, N-1$. $f_j = j/N\tau$ are frequencies, and we may rewrite (10) and (11) into

$$G_j = G(f_j) = \sum_{k=0}^{N-1} e^{-2\pi i f_j t_k} g_k \quad (14)$$

and

$$g_k = g(t_k) = \frac{1}{N} \sum_{j=0}^{N-1} e^{2\pi i f_j t_k} G_j \quad (15)$$

Let us simulate a very noisy cosine signal. We set

```

1  fbar=50;
2  tau=0.001;
3  N=1024;
4  t=tau*[0:N-1];
5  R=2.0;
6  g=cos(2*pi*fbar*t)+R*randn(size(t));
7  plot(t,g,'.');
8  axis([min(t),max(t),-4,4]);
9  print -deps2 ncos1.eps;

```




Figure 5: Noisy cosine (50 Hz) vs. time sampled at millisecond steps

Would you recognize the cosine signal in Fig. 5?

We continue by invoking the fast Fourier transform `fft`:

```

10  G=fft(g);
11  f=[0:N/2-1]/N/tau;
12  S=abs(G(1:N/2)).^2;
13  plot(f,S);
14  print -deps2 ncos2.eps;

```

In Fig. 6 we have plotted the spectral power $S = |G(f)|^2$ for positive frequencies⁷. The prominent peak at $f = 50$ Hz is evident. The remaining spectral power is more or less constant which means white noise.

⁷The spectral power is an even function of f




Figure 6: Spectral power of noisy cosine vs. frequency (Hz)

How can one extract the signal so convincingly from a lot of noise? Not by looking at the sampled data. They appear to be random. However, by performing a Fourier analysis, the different harmonic contributions are excited with their proper phases, so that they add up.

If you know that the signal is spoiled by white noise, you may remove it to a large extent. You might say

```
>> H=G.*(abs(G)>150);
>> h=ifft(H);
```

`ifft` denotes the inverse fast Fourier transform as described by (13).

You can do a lot more with the fast Fourier transform. Here are some examples:

- Differentiation: Fourier transform the signal and multiply by $2\pi if$, and back Fourier transform. Thereby a much better precision can be achieved than by difference quotients.
- Deconvolution: Often the output $b(t) = \int ds r(s)a(t-s)$ is the convolution of a signal a and a transmission function r . If the transmission function is known, then, from $B(f) = R(f)A(f)$, the signal can be reconstructed. Just divide the Fourier transformed output by the Fourier transformed transmission function, and back Fourier transform.

- Differential equations with constant coefficients: Differentiation operators become multiplication operators after Fourier transformation, and differential equations become algebraic equations.
- Data filtering and smoothing: remove the high frequency components.

The fast Fourier transform is an algorithm which takes into account that an operation on $2N$ data points are two operations on N data points. This reduces complexity from N^2 to $N \log(N)$ which makes all the difference.

7 Fitting data

Measured data are to be compared with a theoretically justified model. Discrepancies may arise for two different reasons: the data are not accurate, or the model is inappropriate. Inaccurate data may arise because of systematic or statistical errors.

Let us first formulate the standard task.

You have a set of data $(y_1, x_1), (y_2, x_2), \dots (y_N, x_N)$ and a model $y = f_p(x)$. The p denote a set of parameters which are to be determined. For this we define the variance

$$v(p) = \frac{1}{N-1} \sum_{i=1}^N (f_p(x_i) - y_i)^2 . \quad (16)$$

The optimal fit is defined by minimal variance,

$$v(\bar{p}) = \min_p v(p) . \quad (17)$$

If the model is a linear function, we speak of linear regression. Likewise, if it is a polynomial of degree 2, the data are fitted by quadratic regression, and so forth.

With MATLAB this is very easy. If you say `p=polyfit(x,y,d)`; a polynomial p of degree d is returned which best fits the data vectors x and y . As mentioned earlier, a polynomial is represented by the vector of its coefficients.

Here is an example:

```

1  x=linspace(0,2,256);
2  y=1-x+0.5*x.^2+0.15*randn(size(x));
3  plot(x,y,'.k');
4  hold on;
5  p=polyfit(x,y,2);

```

```

6 yy=polyval(p,x);
7 plot(x,yy,'-k','LineWidth',2);
8 hold off;

```




Figure 7: Noisy quadratic relationship between x (abscissa) and y (ordinate) and reconstruction by quadratic regression

We have plotted the convincing result in Fig. 7.

If the model is not a polynomial, one has to minimize numerically the variance or any other measure of the misfit.

Let us assume a Gaussian peak on top of background. The model is

$$y = y_0 + s e^{-a(x - x_0)^2} . \quad (18)$$

We have four parameters to fit, namely $p = (y_0, s, a, x_0)$. We shall work out an example.

```

1 peak=@(p,x) p(1)+p(2)*exp(-p(3)*(x-p(4)).^2);
2 misfit=@(p,x,y) norm(y-peak(p,x));

```

The first command defines the model, a Gaussian peak. It is the MATLAB equivalent of (18). The next line describes the misfit of a data set \mathbf{x}, \mathbf{y} to the model, for a certain parameter set \mathbf{p} . We will next simulated data:

```

3 tp=[3;1;4;2.5];

```

```

4  x=linspace(0,5,1024);
5  ty=peak(tp,x);
6  ny=ty+0.3*randn(size(x));

```

The true parameter set `tp` gives rise to `ty=peak(tp,x)`. We add noise to it and obtain the noisy data values `ny`. From these noisy data we reconstruct the best fitting parameters `fp` by saying

```

7  fp=fminsearch(misfit,tp,[],x,ny);

```

The first argument to `fminsearch` is the misfit, or cost function. The second argument is a parameter set from which to begin the search. Then comes an options structure for which we specify nothing, i. e. `[]`. In this case the default values for `fminsearch` are used.

Normally the cost function has just one argument, namely the vector of variables over which to minimize. If the cost function requires more arguments (which remain constant when searching for a minimum), these are to be specified in the remainder of the argument list of `fminsearch`. In our case it is the data set `x,ny` to be fitted.

Let us plot the result:

```

8  fy=peak(fp,x);
9  plot(x,ty,'-k',x,ny,'.b',x,fy,'-r','LineWidth',2);

```

However, what will be the result if the search for the best fit starts from another parameter set which is not so close?

Let us try `sp=[1 1 1 1]` as the starting parameter set. This will work. However, with `sp=[0 0 0 0]` the minimization procedure does not converge to a sensible fit.

We might try to fiddle with the optimization options:

```

10  opt=optimset('TolFun',1e-8,'TolX',1e-8,...
11             'MaxFunEvals',20000,'MaxIter',10000);
12  fp=fminsearch(misfit,[0;0;0;0],opt,x,ny);

```

This helps occasionally⁸, but more often not. Obviously the minimization procedure runs into this or another shallow side minimum and misses the absolute minimum.

You might program a coarse pre-search such as this:

⁸Bear in mind that we solve a different problem each time because of randoms in the data set.




Figure 8: A Gaussian peak on top of background: signal (black), signal plus noise (blue dots) and reconstructed signal (red)

```

1  mm=Inf;
2  for j=1:1000
3      p=[6*rand;2*rand;8*rand;5*rand];
4      m=misfit(p,x,ny);
5      if (m<mm)
6          mm=m;
7          pp=p;
8      end;
9  end;

```

It should deliver a good starting parameter set `pp`. The cuboid $0 \leq p_1 \leq 6$, $0 \leq p_2 \leq 2$ etc. has been specified after inspecting the data cloud.

Let us rewrite the above code by using the concepts of anonymous functions, sub-functions, and structures. There is only one function, namely `example_fit` which returns a fit descriptor `fd`. This is a structure which records the details of the fit. We will not comment it line by line because we did so before.

```

1  function fd=example_fit
2  % usage: fd=example_fit
3  % return fit descriptor fd
4  fd.source=mfilename('fullpath');
5  fd.model=@(p,x) p(1)+p(2)*exp(-p(3)*(x-p(4)).^2);

```

```

6  fd.misfit=@(p,x,y) norm(y-fd.model(p,x));
7  fd=simulate_data(fd,1024,0,5,[3 1 4 2.5],0.3);
8  fd=pre_search(fd,1000);
9  fd.fitted_par=fminsearch(fd.cost_fun,fd.start_par);
10 plot_fit(fd);
11 end
12
13 function fd=simulate_data(fd,nd,xlo,xhi,tp,noise)
14 fd.xdata=linspace(xlo,xhi,nd);
15 fd.true_par=tp;
16 fd.noise=noise;
17 fd.ydata=fd.model(tp,fd.xdata)+noise*randn(size(fd.xdata));
18 fd.cost_fun=@(p) fd.misfit(p,fd.xdata,fd.ydata);
19 end
20
21 function fd=pre_search(fd,n)
22 fd.presearches=n;
23 bmf=Inf; % initial best misfit
24 for k=1:n
25     rp=[6*rand 2*rand 8*rand 5*rand]; % random parameters
26     rmf=fd.cost_fun(rp); % random mistfit
27     if rmf<bmf
28         bmf=rmf;
29         bp=rp; % current best parameters
30     end;
31 end
32 fd.start_par=bp;
33 end
34
35 function plot_fit(fd)
36 x=fd.xdata;
37 yt=fd.model(fd.true_par,x); % true
38 ys=fd.ydata; % simulated
39 yf=fd.model(fd.fitted_par,x); % fitted
40 plot(x,ys,'.b',x,yt,'-k',x,yf,'-r','LineWidth',3);
41 end
42
43 % eof example_fit

```

The fit of simulated data to the model is the same as before. The interaction with the outside world consists now in only one .m file in the file system and only one variable fd in the workspace. By this we minimize clashes of filenames and variables.

If you say fd on the command line you will see something like this:

```

1      source: [1x56 char]
2      model: @(p,x)p(1)+p(2)*exp(-p(3)*(x-p(4)).^2)
3      misfit: @(p,x,y)norm(y-fd.model(p,x))
4      xdata: [1x1024 double]
5      true_par: [3 1 4 2.5000]
6      noise: 0.3000
7      ydata: [1x1024 double]
8      cost_fun: @(p)fd.misfit(p,fd.xdata,fd.ydata)
9      presearches: 1000
10     start_par: [3.0837 0.3074 4.8080 2.6328]
11     fitted_par: [3.0080 0.9388 3.7287 2.4874]

```

`source` is the path to the program which has produced this output. `cost_fun` is defined with the currently available `xdata` and `ydata`. They are frozen for this purpose. The rest explains itself.

8 Simulated Annealing

Minimization problems are among the most frequently encountered tasks of computational physics. There is a set of parameters, and each parameter set is weighted by a cost function. The problem to be solved is: find the parameter set for which the costs are minimal.

We did already discuss such an optimization problem in the context of model fitting. The Nelder-Mead simplex method (`fminsearch`) is never the best, but always a good choice to tackle the problem. Remember: you have to specify a starting point when searching for the minimum. In general, the local minimum closest to the starting point is found.

If there are very many parameters, there are also very many local minima, but there is only one global minimum. Therefore, a coarse search for a sensible starting point is the first step in order to find a sensible starting point. In many cases this first step is visual inspection of data, or intuition, or the result of previous minimization efforts.

There are however problems where there are so many local minima that a direct search for the global minimum is appropriate.

Here we discuss a computational classic: the traveling salesman problem.

There are NC cities to be visited, their Cartesian coordinates are stored in vectors `XC` and `YC`, respectively. The itinerary is described by a permutation of indices from 1 to NC , a vector `it`. The cost function is the length of a round trip, as calculated by

```

1  % this file is ts_length.m
2  function len=ts_length(it);

```

```

3  global NC XC YC
4  itt=it([2:NC,1]);
5  len=sum(sqrt((XC(itt)-XC(it)).^2+(YC(itt)-YC(it)).^2));

```

The `global` statement says that all or some of the variables may be visible to subprograms or functions if they want to see them.

We silently assume that the effort to travel between two cities is proportional to their distance. This can be easily modified by introducing a table which describes the effort to travel from one to another city.

The number of possible itineraries is finite, but prohibitively large. If there are only 20 cities, we have to check $19! \approx 10^{17}$ possibilities. If one check lasts only 1 μ s, this requires 10^{11} s, or 4000 years.

Here we describe the simulated annealing algorithm⁹. A probe is heated and then slowly cooled down. Thereby all sorts of defects may be mended. Temperature allows for random fluctuations, and cooling will lead to the state of lowest energy. Even if the system is close to a local energy minimum, a fluctuation may send it to an even better minimum.

If a new configurations has a lower energy (cost), it is accepted straight away. If its energy is higher, it shall be accepted with probability $\exp(-\Delta E/T)$ where T is the temperature and ΔE the energy increase.

So we set up a sensible temperature, allow for random variations of the itinerary, and let the system cool down.

Here is a not-yet optimized algorithm.

```

1  % this file is ts_problem.m
2  global NC XC YC
3  NC=20;
4  XC=rand(1,NC);
5  YC=rand(1,NC);
6  T=ts_hot;
7  suc=1;
8  it=[1:NC];
9  while suc>0
10     [it,suc]=ts_anneal(T,it);
11     T=0.9*T;
12 end

```

We simulate the locations of 20 cities, set up an initial temperature T , choose city 1 to city 2 to ... as an initial itinerary, and let it cool down in steps of 10 % temperature reduction. For each temperature, the system is annealed.

⁹Adapted from W. H. Press, B. P. Flannery, S. A. Teukolsky and W. T. Vetterling, Numerical Recipes, Cambridge University Press 1986, ISBN 0521308119

The relevant function returns the new best itinerary and the number of successful trials to find a shorter path. The program stops if there are no more improvements.




Figure 9: The shortest itinerary found by the simulated annealing algorithm

Fig. 9 shows the solution. It was created by calling `ts_problem` (see above). The `anneal` function lets the system perform random fluctuations. If a modified itinerary is shorter, it is accepted right away. If it is longer, but not too much, it will be accepted as well. How much is too much is ruled by the temperature. The lower the temperature, the likelier a worsening will be rejected.

Here is the code:

```
1 % this file is ts_anneal.m
2 function [ii,suc]=ts_anneal(T,it);
3 global NC
4 max_rep=100*NC;
5 max_suc=10*NC;
6 suc=0;
7 ii=it;
8 min_len=ts_length(it);
```

```

9   for rep=1:max_rep
10      im=ts_modify(it);
11      len=ts_length(im);
12      if rand<exp(-(len-min_len)/T)
13          ii=im;
14      end;
15      if len<min_len
16          min_len=len;
17          suc=suc+1;
18      end
19      if suc==max_suc
20          break
21      end
22  end

```

We try at most `max_rep` times. And if the number of successes is too large, we stop as well and try a lower temperature. The decisive statement is

```
>> if rand<exp(-(len-min_len)/T)
```

expressing that we accept a modification if the new length is smaller than the up-to-now lowest length, or if the new length is not too much larger. This is not yet, but might become an improvement because the modification may be close to a better local minimum.

It remains to be explained how itineraries are modified at random. We have programmed two alternatives. We cut out a sub-itinerary and revert it, with 50 % probability, or we transport it by a random amount. This is formulated here:

```

1   % this file is ts_modify.m
2   function im=ts_modify(it);
3   global NC
4   n=sort(ceil(NC*rand(1,2)));
5   ib=it(1:n(1)-1);
6   is=it(n(1):n(2));
7   ia=it(n(2)+1:NC);
8   if rand<0.5
9       im=[ib,flip1r(is),ia];
10  else
11      ir=[ib,ia];
12      len=length(ir);
13      n=ceil(len*rand);
14      im=[ir(1:n),is,ir(n+1:len)];
15  end;

```

The itinerary `it` is split into a randomly chosen sub-itinerary `is`, its leader `ib` and trailer `ia`. `fliplr` flips matrices in left-right direction. Note that this simple code does not guarantee that the original and the modified itinerary differ.

In order to be complete, we also show how the initial temperature was chosen:

```

1  % this file is ts_hot.m
2  function T=ts_hot();
3  global NC
4  lmax=0;
5  for k=1:20
6      len=ts_length(randperm(NC));
7      if len>lmax
8          lmax=len;
9      end;
10 end;
11 T=2*lmax;

```

Calling `randperm(N)` returns a random permutation of `[1:N]`.

9 Finite difference method

The finite difference method of treating differential equations is simple. The limit in

$$f'(x) = \lim_{h \rightarrow 0} \frac{f(x + h/2) - f(x - h/2)}{h} \quad (19)$$

is approximated by a finite difference quotient. Likewise, the second derivative is to be approximated by

$$f''(x) \approx \frac{f(x + h) - 2f(x) + f(x - h)}{h^2} . \quad (20)$$

The Laplacian in two dimensions is approximated by

$$\frac{f(x + h, y) + f(x, y + h) + f(x - h, y) + f(x, y - h) - 4f(x, y)}{h^2} . \quad (21)$$

Generalizations to different step widths h_x, h_y and for more than two dimensions are obvious.

Our first example is a simple ordinary differential equation:

$$f'' = -f \text{ with } f(0) = 0 \text{ and } f(\pi/2) = 1 , \quad (22)$$

the solution of which is $f(x) = \sin(x)$.

The normal ordinary differential equation solvers require initial conditions, such as $f(0) = 0$ and $f'(0) = 1$. We, however, are faced with boundary conditions.

The following program realizes the finite difference method.

```
1  % this file is fdm_sin.m
2  N=16;
3  h=pi/2/N;
4  next=(1/h^2)*ones(1,N-2);
5  main=(-2/h^2+1)*ones(1,N-1);
6  DE=diag(next,-1)+diag(main,0)+diag(next,1);
7  RS=zeros(1,N-1)';
8  RS(N-1)=-1/h^2;
9  sol=DE\RS;
10 x=[h:h:(N-1)*h];
11 xx=linspace(0,pi/2,256);
12 plot(x,sol,'r.',xx,sin(xx),'MarkerSize',20);
13 axis tight;
14 print -depsc fdm_sin.eps;
```




Figure 10: Analytic solution (full line) and finite difference approximation (full circles) of $f'' + f = 0$ with $f(0) = 0$ and $f(\pi/2) = 1$

We have plotted the result in Fig. 10. The maximum deviation is 2×10^{-4} although only 15 representative points have been used. Well, better say

17 because the entries $RS(1)=0$ and $RS(N-1)=-1/h^2$ reflect the boundary conditions. The second derivatives at the smallest and largest interior point require information from outside.

Note the `diag` function which allows to set and to extract diagonals. Also note that `x=M\y` is short for solving the system of linear equations $y=M*x$. `axis tight` does what it says: the axis frame is fitted as tightly as possible to the data.

In our second example we want to solve a two-dimensional boundary problem. Let us reproduce the MATLAB logo which is a visualization of the lowest order eigenmode of the following equation:

$$-\Delta u = \Lambda u \quad . \quad (23)$$

$u = u(x, y)$ is defined on a set $\Omega \subset \mathbb{R}^2$ with appropriate conditions on the boundary $\partial\Omega$. It is a challenge to solve this partial differential eigenvalue problem on the L-shaped domain

$$\Omega = [-1, 1] \times [-1, 1] - [-1, 0] \times [-1, 0] \quad (24)$$

with $u = 0$ on $\partial\Omega$.

We shall not describe the most efficient solution but a program which is easy to generalize.

We begin by describing the domain.

```

1  % this file is ml_logo.m
2  N=32;
3  h=2/(N-1);
4  x=linspace(-1,1,N);
5  y=linspace(-1,1,N);
6  [X,Y]=meshgrid(x,y);
7  Domain=(abs(X)<1)&(abs(Y)<1)&((X>0)|(Y>0));
```

X and Y are $N \times N$ matrices such that $X(j,k)=x(k)$ and $Y(j,k)=y(j)$. `Domain` is likewise an $N \times N$ matrix with 1 for an interior point and 0 for a boundary or exterior point.

A subprogram `laplace` will calculate the sparse matrix L approximating the Laplacian Δ by (21). It also returns the mapping of the running index `a` for variables to pairs j, k which index the x, y coordinates. L is defined by $(\Delta u)_a = \sum_b L_{ab} u_b$. We have $L_{aa} = -4/h^2$ and $L_{ab} = 1/h^2$ whenever a and b index neighboring variables.

The rest is easy. We call the sparse matrix eigenvalue solver `eigs` which returns the lowest order eigensolution u and the corresponding eigenvalue d :

```

8 [L,J,K]=laplace(Domain,h);
9 [u,d]=eigs(-L,1,'sm');
10 s=sign(sum(u));
11 field=zeros(size(Domain));
12 for a=1:size(u)
13     field(J(a),K(a))=s*u(a);
14 end;
15 mesh(field);
16 axis off;
17 print -depsc ml_logo.eps;

```




Figure 11: The lowest order eigenmode of $-\Delta u = \Lambda u$ on an L-shaped domain has become their logo

The `laplace` function which calculates the Laplacian as a sparse matrix and returns the indexing scheme is here:

```

1 % this file is laplace.m
2 function [L,J,K]=laplace(D,h);
3 [Nx,Ny]=size(D);
4 jj=zeros(Nx*Ny,1);
5 kk=zeros(Nx*Ny,1);
6 aa=zeros(Nx,Ny);
7 Nv=0;
8 for j=1:Nx
9     for k=1:Ny

```

```

10     if D(j,k)==1
11         Nv=Nv+1;
12         jj(Nv)=j;
13         kk(Nv)=k;
14         aa(j,k)=Nv;
15     end
16 end
17 end
18 L=sparse(Nv,Nv);
19 for a=1:Nv
20     j=jj(a);
21     k=kk(a);
22     L(a,a)=-4/h^2;
23     if D(j+1,k)==1
24         L(a,aa(j+1,k))=1/h^2;
25     end
26     if D(j-1,k)==1
27         L(a,aa(j-1,k))=1/h^2;
28     end
29     if D(j,k+1)==1
30         L(a,aa(j,k+1))=1/h^2;
31     end
32     if D(j,k-1)==1
33         L(a,aa(j,k-1))=1/h^2;
34     end
35 end
36 J=jj(1:Nv);
37 K=kk(1:Nv);

```

In our case, we have 675 independent variables, hence L should require 3.6 MB of storage. However, most of its entries are zeros. Therefore it is sufficient to store the locations and values of the non-vanishing entries only, 42 kB in this case. MATLAB provides this sparse matrix technology which is essential for partial differential equations and for many other branches of computational physics.

10 Propagation

In the previous section we have studied a special class of partial differential equations. The solution was fixed by values on a boundary. Here we address another class of problems. A field is prescribed at time $t = 0$. How will it propagate in the course of time?

Let us study a simple, but typical¹⁰ partial differential equation of this type, namely

$$\dot{u} = \Delta u \quad , \quad (25)$$

for simplicity in one spatial dimension, $u = u(t, x)$. The dot denotes partial differentiation with respect to the time argument t , the Laplacian is the second derivative with respect to the space argument x . The spatial region is an interval, $0 \leq x \leq 1$, say. Time begins with $t = 0$. We look for a solution of (25) subject to the following conditions:

$$u(0, x) = u_0(x) \quad \text{as well as} \quad u(t, 0) = b_0(t) \quad \text{and} \quad u(t, 1) = b_1(t) \quad . \quad (26)$$

$u_0 = u_0(x)$ is an initial condition, $b_k = b_k(t)$ for $k = 0, 1$ are boundary conditions. We assume that these restrictions are compatible, i. e. $u_0(0) = b_0(0)$ and $u_0(1) = b_1(0)$.

We resort to the finite difference method and write

$$u_{j,k} = u(j\tau, kh) \quad (27)$$

for integers j and k and time and space steps τ and h , respectively.

The most natural approach is to approximate (25) by

$$u_{j+1,k} = u_{j,k} + \tau \left\{ \frac{u_{j,k+1} - 2u_{j,k} + u_{j,k-1}}{h^2} \right\} \quad . \quad (28)$$

If the variables are labeled by $k = 1, 2, \dots, N$, then $u_{j,0}$ in (28) has to be replaced by $b_{0,j} = b_0(j\tau)$ and $u_{j,N}$ by $b_{1,j} = b_1(j\tau)$. The differencing scheme is fully explicite since it allows to calculate new field values in terms of previously known values. One may write

$$u_{j+1} = (I + \tau L) u_j \quad . \quad (29)$$

It turns out that (28) is stable¹¹ only if $\tau < h^2$. Hence, with $h = 0.01$, we must proceed in time steps of $\tau = 10^{-4}$ or smaller. This is not acceptable.

Another, fully implicite differencing scheme is

$$u_{j,k} = u_{j+1,k} - \tau \left\{ \frac{u_{j+1,k+1} - 2u_{j+1,k} + u_{j+1,k-1}}{h^2} \right\} \quad , \quad (30)$$

or

$$u_j = (I - \tau L) u_{j+1} \quad (31)$$

¹⁰heat transport, diffusion, etc.

¹¹see the section on *Diffusive Initial Value Problems in Numerical Recipes*, loc. cit.

which is solved by

$$u_{j+1} = (I - \tau L)^{-1} u_j . \quad (32)$$

This one is stable for all time steps τ , but we have to solve a system of linear equations for every propagation step. Both differencing schemes are biased: they approximate, by $\frac{f(t + \tau) - f(t)}{\tau}$, the time derivative at t or at $t + \tau$, respectively, while it should be at $t + \tau/2$.

We therefore write

$$u_{j+1/2} = (I + \frac{\tau L}{2}) u_j = (I - \frac{\tau L}{2}) u_{j+1} \quad (33)$$

which amounts to

$$u_{j+1} = (I - \frac{\tau L}{2})^{-1} (I + \frac{\tau L}{2}) u_j . \quad (34)$$

This so called *Crank-Nicholson scheme* is stable for all time steps τ and one order more accurate than (29) or (31).

The following MATLAB function propagates the field u at time t by one time step τ according to the Crank-Nicholson scheme. The spacing h and the boundary values $b_k(t + \tau/2)$ (a two component vector) have to be specified as well¹². Fields are to be represented as column vectors.

```

1  % this file is cn_step.m
2  function v=cn_step(u,tau,h,b);
3  N=length(u);
4  z=0.5*tau/h^2;
5  L=laplacian(N);
6  vv=u+z*L*u;
7  vv(1)=vv(1)+2*z*b(1);
8  vv(N)=vv(N)+2*z*b(2);
9  M=eye(N)-z*L;
10 v=M\vv;
11
12 function L=laplacian(N);
13 next=ones(1,N-1);
14 main=ones(1,N);
15 L=diag(next,-1)-2*diag(main,0)+diag(next,1);

```

Let us study a situation where $u_0(x) = \sin \pi x + \sin 2\pi x$, $b_0(t) = 0$ and $b_1(t) = 0$.

¹²more precisely: the mean of the boundary values before and after the time step

```

1  % this file is cn_example.m
2  Nx=500;
3  x=linspace(0,1,Nx)';
4  h=x(2)-x(1);
5  tau=0.0025;
6  u0=sin(pi*x)+sin(2*pi*x);
7  u=u0(2:Nx-1);
8  b=[0,0];
9  field=u;
10 for j=1:100
11     u=cn_step(u,tau,h,b);
12     field=[field,u];
13 end;
14 contour(field,32);
15 print -depsc cn_example.eps;

```




Figure 12: Contour plot of $u = u(t, x)$. Crank-Nicolson scheme, 100 time propagation steps, t running from left to right. Initially, $u(0, x) = \sin(\pi x) + \sin(2\pi x)$

By the way, Jean Baptiste Joseph Fourier, who lived from 1768 to 1830, has invented the method of decomposing functions into harmonic contributions while trying to solve the heat equation (25). In today's notation:

$$u(t, x) = \sum_{n=1}^{\infty} c_n(t) \sin n\pi x \quad (35)$$

is solved by

$$c_n(t) = c_n(0) e^{-(n\pi)^2 t} . \quad (36)$$

Thus, the $\sin(2\pi x)$ contribution dies much more rapidly than the $\sin(\pi x)$ term as can be read off Fig. 12.

A Figures Makefile

The figures (as .pdf-files) were produced with the aid of .m-files printed in this document and by the following function. Note that '!...' invokes the system command shell.

```
1  function npfigs()
2  % usage: npfigs
3  % produces *.pdf figures for np.tex
4
5  disp 'make planck.pdf'
6  x=linspace(0,10,512);
7  s=planck(x);
8  plot(x,s,'-k');
9  print -deps2 planck.eps;
10 eval('!epstopdf planck.eps');
11 clear all;
12
13 disp 'make kepler1.pdf kepler2.pdf kepler3.pdf'
14 [t,y]=ode45(@kepler,0:0.1:50,[1;0;0;0.8]);
15 plot(y(:,1),y(:,3),0,0,'+k');
16 axis square;
17 print -deps2 kepler1.eps;
18 eval('!epstopdf kepler1.eps');
19 Ekin=0.5*(y(:,2).*y(:,2)+y(:,4).*y(:,4));
20 Epot=-1./sqrt(y(:,1).*y(:,1)+y(:,3).*y(:,3));
21 plot(t,Ekin+Epot);
22 print -deps2 kepler2.eps;
23 eval('!epstopdf kepler2.eps');
24 tol=odeset('RelTol',1e-8);
25 [~,y]=ode45(@kepler,0:0.1:50,[1;0;0;0.8],tol);
26 plot(y(:,1),y(:,3),0,0,'+k');
27 axis square;
28 print -deps2 kepler3.eps;
29 eval('!epstopdf kepler3.eps');
30 clear all;
31
32 disp 'make ncos1.pdf ncos2.pdf'
33 reset(RandStream.getDefaultStream);
34 noisy_cos;
35 eval('!epstopdf ncos1.eps');
36 eval('!epstopdf ncos2.eps');
37 clear all;
38
```

```

39 disp 'make fitr.pdf fitp.pdf'
40 reset(RandStream.getDefaultStream);
41 fit_regr;
42 eval('!epstopdf fitr.eps');
43 clear all;
44 fit_peak;
45 print -depsc fitp.eps
46 eval('!epstopdf fitp.eps');
47 clear all;
48
49 disp 'make ts_problem.pdf'
50 reset(RandStream.getDefaultStream);
51 ts_problem;
52 plot(XC(it),YC(it),'r',[XC(it),XC(it(1))],[YC(it),YC(it(1))],...
53      '-k','MarkerSize',20);
54 axis equal;
55 axis off;
56 print -depsc ts_problem.eps;
57 eval('!epstopdf ts_problem.eps');
58 clear all;
59
60 disp 'make fdm_sin.pdf'
61 fdm_sin;
62 eval('!epstopdf fdm_sin.eps');
63 clear all
64
65 disp 'make ml_logo.pdf'
66 ml_logo;
67 eval('!epstopdf ml_logo.eps');
68 clear all;
69
70 disp 'make cn_example.pdf'
71 cn_example;
72 eval('!epstopdf cn_example.eps');
73 clear all;
74
75 delete *.eps

```

List of Figures

1	The black body radiation spectral intensity $S(x)$ where $x = \hbar\omega/k_{\text{B}}T$	9
2	Planetary motion without explicit accuracy control	13
3	Total energy vs. time without accuracy control	13
4	Planetary motion with accuracy control	14
5	Noisy cosine (50 Hz) vs. time sampled at millisecond steps	16
6	Spectral power of noisy cosine vs. frequency (Hz)	17
7	Noisy quadratic relationship between x (abscissa) and y (ordinate) and reconstruction by quadratic regression	19
8	A Gaussian peak on top of background: signal (black), signal plus noise (blue dots) and reconstructed signal (red)	21
9	The shortest itinerary found by the simulated annealing algorithm	25
10	Analytic solution (full line) and finite difference approximation (full circles) of $f'' + f = 0$ with $f(0) = 0$ and $f(\pi/2) = 1$	28
11	The lowest order eigenmode of $-\Delta u = \Lambda u$ on an L-shaped domain has become their logo	30
12	Contour plot of $u = u(t, x)$. Cranck-Nicolson scheme, 100 time propagation steps, t running from left to right. Initially, $u(0, x) =$ $\sin(\pi x) + \sin(2\pi x)$	34